

VISUALIZING HYPERBOLIC HONEYCOMBS

ROICE NELSON AND HENRY SEGERMAN

ABSTRACT. We explore visual representations of tilings corresponding to Schläfli symbols. In three-dimensions, we call these tilings “honeycombs”. Schläfli symbols encode, in a very efficient way, regular tilings of spherical, euclidean and hyperbolic spaces in all dimensions. In three-dimensions, there are only a finite number of spherical and euclidean honeycombs, but infinitely many hyperbolic honeycombs. Moreover, there are only four hyperbolic honeycombs with material vertices and material cells (the cells are entirely inside of hyperbolic space), eleven with ideal vertices or cells (the cells touch the boundary of hyperbolic space in some way), and all others have either hyperideal vertices or hyperideal cells (the cells go outside of the boundary of hyperbolic space in some way). We develop strategies for visualising honeycombs in all of these categories, either via rendered images or 3D prints. High resolution images are available at hyperbolichoneycombs.org.

“Art attempts to portray in limited space, and song in time, what is otherwise unlimited.” [4]

1. INTRODUCTION

Schläfli symbols are an extremely efficient system of notation for regular polytopes, euclidean and hyperbolic tilings. They are defined recursively, as follows. A Schläfli symbol of the form $\{p\}$ denotes the regular p -gon. A Schläfli symbol of the form $\{p, q\}$ denotes a two-dimensional tiling of p -gons in which q such polygons meet at each vertex. For example, $\{4, 3\}$ denotes the tiling of squares, with three arranged around each vertex, also known as the cube. A Schläfli symbol of the form $\{p, q, r\}$ denotes a three dimensional tiling of *cells* with Schläfli symbol $\{p, q\}$, with r such cells meeting at each edge. For example, $\{4, 3, 4\}$ denotes the tiling of three-dimensional euclidean space by cubes. In three-dimensions, we also call these tilings *honeycombs*.

In Section 2 we warm up by looking at two-dimensional tilings corresponding to Schläfli symbols of the form $\{p, q\}$, and view these in spherical, euclidean and hyperbolic spaces. We also note the relation between Schläfli symbols and duality, and the rule to determine which geometry a tiling lives in from its symbol. In Section 3, we move on to three-dimensional honeycombs, again noting the relation with duality and the rule for the geometry of a honeycomb. In Section 4, we introduce tilings and honeycombs with ideal vertices and ideal cells. In Section 5 we give a visual map of the different kinds of geometries and vertex/cell types of honeycombs. In Section 6 we introduce tilings and honeycombs with hyperideal vertices and hyperideal cells. In Section 7, we draw pictures (in various forms) of the honeycombs on the boundary of the infinite cube of three-dimensional Schläfli symbols, and on surfaces within the cube corresponding to

honeycombs with ideal vertices or ideal cells. In Section 8, we display artworks, both 3D printed sculpture and 2D renders, based on our investigations. In Section 9 we give some details on how we created our images and models, and in Section 10 we give some directions for future work.

2. TWO-DIMENSIONS

2.1. Geometries. Let's start by looking at some examples of two-dimensional tilings given by Schläfli symbols, see Figure 1. Immediately there is a question of how to draw pictures of these tilings. The $\{4, 4\}$ tiling is very familiar, and can be drawn on the page without any distortion, but the other two tilings cannot. The $\{4, 3\}$ tiling is otherwise known as the cube: three squares meet around each vertex. However, for our purposes it is better to tile the sphere rather than the cube, since then we are tiling an object with uniform geometry – i.e. a surface with constant curvature. We can get from the polyhedral version of the cube to the tiling of the sphere by radially projecting the vertices, edges and faces of a cube inscribed in the sphere to the sphere. Figure 1a shows the stereographic projection [8, pp. 140-148] of the $\{4, 3\}$ tiling of the sphere to the plane. In Figure 1c we draw the $\{4, 5\}$ tiling on the Poincaré disk model of \mathbb{H}^2 [8, pp. 315-319], which is again a space with constant curvature, this time negative curvature.




FIGURE 1. Two dimensional tilings by squares.

2.2. Duality. As is well-known, the regular polyhedra relate to each other by duality: if we plot a vertex at the center of each face of a cube, and connect these new vertices by edges that cross the edges of the cube, then we get an octahedron. If we start with an octahedron and perform the same operations, we get back the cube. Similarly the dodecahedron and icosahedron are duals of each other, and the tetrahedron is self-dual. In terms of the Schläfli symbols, we get the dual of a polyhedron by reversing its Schläfli symbol. This relation extends from polyhedra (tilings of the sphere) to tilings of the euclidean plane ($\{3, 6\}$ is dual to $\{6, 3\}$ and $\{4, 4\}$ is self-dual) and tilings of the hyperbolic plane.




FIGURE 2. Dual two-dimensional tilings.

There is more than one way to draw a tiling of hyperbolic space – in particular Figure 1c shows a *cell-centered* picture of the $\{4, 5\}$ tiling, and Figure 2c shows a *vertex-centered* picture of the $\{5, 4\}$ tiling, so called because a cell (respectively, vertex) of the $\{4, 5\}$ (respectively, $\{5, 4\}$) tiling is at the center of the Poincaré disk.

2.3. Which two-dimensional Schläfli symbols produce tilings of which geometries? Suppose we try to tile the euclidean plane with polygons as given to us by a Schläfli symbol $\{p, q\}$. Then q p -gons meet at a vertex, each of which has internal angle $(1 - 2/p)\pi$, for a total angle of $q(1 - 2/p)\pi$. We will be able to tile the euclidean plane if this value is equal to 2π , or equivalently if

$$(1) \quad (p - 2)(q - 2) = 4$$

If the expression on the left hand side is less than 4, then we instead get a polyhedron [3, p. 5], which as before we view as a tiling of the sphere. If this value is greater than 4, then we can tile the hyperbolic plane [2, pp. 155-156]. See Figure 3.




FIGURE 3. Geometries associated to the two-dimensional Schläfli symbols: below the curve is spherical, on the curve is euclidean, above the curve is hyperbolic.

3. THREE-DIMENSIONS

3.1. Geometries. Three-dimensional Schläfli symbols also represent tilings of various spaces with constant curvature, analogous to the sphere, S^2 , the euclidean plane, \mathbb{E}^2 , and the hyperbolic plane, \mathbb{H}^2 , in two-dimensions. Here, we tile the three-sphere, S^3 , three-dimensional euclidean space, \mathbb{E}^3 , and three-dimensional hyperbolic space, \mathbb{H}^3 . In three-dimensions, we also call these tilings *honeycombs*. Figure 4 shows spherical, euclidean and hyperbolic honeycombs made out of cubes.




FIGURE 4. Three-dimensional honeycombs with cubical cells.

Figure 4a shows (the edges of) the tiling of S^3 by eight cubes. As a four-dimensional polytope, this is the hypercube – three cubes meet around each edge, corresponding

to the last term of its Schläfli symbol $\{4, 3, 3\}$. As with the two-dimensional tilings of S^2 we can get from the four-dimensional polytope to the tiling (honeycomb) of S^3 by radially projecting the polytope onto S^3 . We draw the hypercube by first doing this radial projection, then stereographically projecting to euclidean space [11]. Figure 4b shows the familiar tiling of \mathbb{E}^3 by cubes. Figure 4c shows the $\{4, 3, 5\}$ tiling of the Poincaré ball model of \mathbb{H}^3 . Or rather, this shows a small part of the tiling – there is a central cube, surrounded by a number of layers of cubes. The full tiling fills up all of \mathbb{H}^3 , so a picture of it would just look like a solid ball. Figure 5 shows three layers of the $\{4, 3, 5\}$ tiling: the central cube, the cubes adjacent to the central cube, and the cubes adjacent to those.




FIGURE 5. Three layers of cubes in the $\{4, 3, 5\}$ honeycomb.

3.2. Duality. As in two-dimensions, the three-dimensional Schläfli symbols also have a notion of duality, and again we can get from a tiling to its dual by reversing the Schläfli symbol. See Figure 6. Geometrically, the process to get the dual of a honeycomb goes as follows: First we put a new vertex at the center of each three-dimensional cell of the initial honeycomb. Next, we join two new vertices by a new edge if the corresponding cells of the old honeycomb share a face. We add a new face cutting through each edge of the old honeycomb, and finally we add a new cell corresponding to each vertex of the old honeycomb.

3.3. Which three-dimensional Schläfli symbols produce honeycombs of which geometries? As in two-dimensions, there is a formula that tells us whether the honeycomb corresponding to the Schläfli symbol $\{p, q, r\}$ tiles S^3 , \mathbb{E}^3 or \mathbb{H}^3 . Instead of adding up the internal angles of polygons around a vertex, we add up the dihedral angles of polyhedra around an edge. For a euclidean polyhedron given by the Schläfli symbol $\{p, q\}$, the dihedral angle is $2 \arcsin \left(\cos \frac{\pi}{q} / \sin \frac{\pi}{p} \right)$, for a total angle of $2r \arcsin \left(\cos \frac{\pi}{q} / \sin \frac{\pi}{p} \right)$ [2, p. 157]. We will be able to tile \mathbb{E}^3 if this value is equal to 2π , or equivalently if




FIGURE 6. Three-dimensional honeycombs dual to the honeycombs in Figure 4. Each image shows eight cells from the honeycomb, and one cubical cell from the corresponding dual honeycomb.

$$(2) \quad \cos \frac{\pi}{q} = \sin \frac{\pi}{p} \sin \frac{\pi}{r}.$$



If the left hand side is less than the right, then we instead get a four-dimensional polytope, which as before we view as a tiling of S^3 , and if the left hand side is greater than the right, then we can tile \mathbb{H}^3 [2, p. 157].

Other than the complexity of drawing the figures, and the complexity of the formula, the story in three-dimensions is so far very similar to the story in two-dimensions. However, some new features arise in three-dimensions, as we will see in the next sections.


4. IDEALIZATION

4.1. Ideal vertices. If we continue the sequence of tilings shown in Figure 1, all Schläfli symbols $\{4, q\}$, $q \geq 5$, are tilings of \mathbb{H}^2 . See Figure 7. For each such tiling, the internal angle of each square must be $2\pi/q$, so as q increases, the square gets spikier, and the vertices move farther and farther away from the center of the square. In the limit, the Schläfli symbol $\{4, \infty\}$ corresponds to a tiling with infinitely many squares around each vertex, each of which has internal angle zero, and the vertices are infinitely far away, on the boundary of \mathbb{H}^2 . See Figure 7d.

A vertex that is infinitely far away, on the boundary of hyperbolic space, is called an *ideal vertex*. We call a vertex that is inside hyperbolic space a *material vertex*. In three-dimensions, we don't need to go to Schläfli symbols with infinite terms to find ideal vertices. In fact, the next Schläfli symbol in the sequence shown in Figure 4, $\{4, 3, 6\}$, corresponds to a honeycomb of \mathbb{H}^3 by cubes with ideal vertices. See Figure 8. At first, it is a little unclear why $\{4, 3, 6\}$ should have ideal vertices while $\{4, 3, 5\}$ does not. This is easiest to see with the dual honeycomb, which we investigate in the next section.

FIGURE 7. Tilings given by Schläfli symbols $\{4, q\}$.FIGURE 8. Three layers of cubes in the $\{4, 3, 6\}$ honeycomb.

4.2. Ideal tiles. If a tiling has ideal vertices, then its dual will have ideal tiles (i.e. ideal faces for two-dimensional tilings, and ideal cells for three-dimensional honeycombs). We can see this in the two-dimensional case by looking at the dual sequence of tilings to the sequence shown in Figure 7. See Figure 9.

FIGURE 9. Tilings given by Schläfli symbols $\{p, 4\}$.

We cannot draw a vertex-centered picture of a tiling with ideal vertices, since the vertices cannot be put at the center of the Poincaré disk. Similarly, it is impossible




FIGURE 10. A single cell from the tiling given by the Schläfli symbol $\{6, 3, 4\}$.

to draw a tile-centered picture of a tiling with ideal tiles. Figure 10 shows one ideal cell from the $\{6, 3, 4\}$ honeycomb. The cell has infinitely many hexagonal faces, as you would expect since its Schläfli symbol is $\{6, 3\}$. Geometrically, the vertices of the cell sit on a *horosphere* [8, p. 326] of \mathbb{H}^3 . The Poincaré ball model represents \mathbb{H}^3 as a euclidean ball in \mathbb{E}^3 . A horosphere is represented in the model as a euclidean sphere which is contained in and touches the boundary of the Poincaré ball. As viewed in the negatively curved metric in \mathbb{H}^3 however, a horosphere is isometric to a copy of the euclidean plane \mathbb{E}^2 . This makes sense, since $\{6, 3\}$ is a tiling of the euclidean plane.

In general, an *ideal cell* is a three-dimensional cell whose vertices lie on a horosphere. We call a cell whose vertices lie on a sphere that is strictly inside of \mathbb{H}^3 a *material cell*.

From this perspective we can see why $\{4, 3, 6\}$ has ideal vertices: the dual honeycomb has cells with Schläfli symbol $\{6, 3\}$. Since this is a tiling of the euclidean plane rather than a sphere, the cell vertices must live on a horosphere. The original honeycomb $\{4, 3, 6\}$ has a vertex at the center of this dual cell, and the center of a horosphere is on the boundary of \mathbb{H}^3 .

Note that only the vertices of an ideal cell live on the horosphere, not the edges or faces. In hyperbolic space, the euclidean geometry of the horosphere curves away from the hyperbolic geodesics of the edges and faces.

5. THE THREE-DIMENSIONAL SCHLÄFLI SYMBOL MAP

We can now describe the map of the kinds of geometry, vertices and cells we see for three-dimensional Schläfli symbols. See Figure 11. Three surfaces are shown. In the bottom left is the surface associated to Equation 2. Below the surface the honeycomb is spherical, on the surface the honeycomb is euclidean, and above the surface the honeycomb is hyperbolic. The other two surfaces are associated to Equation 1, the

first using p and q from the three-dimensional Schläfli symbol $\{p, q, r\}$, and the second using r and q . Honeycombs on the first surface have ideal cells, while those closer to the origin have material cells. Honeycombs on the second surface have ideal vertices, while those closer to the origin have material vertices. The three honeycombs lying on both of these surfaces have ideal cells and ideal vertices.




FIGURE 11. Geometries and cell/vertex types associated to the three-dimensional Schläfli symbols.

In the next section, we explore what happens when we move to the far side of these two surfaces.

6. HYPERIDEALIZATION

6.1. Hyperideal vertices. Let's look at what happens when we make the $\{4, \infty\}$ squares of Figure 7 even bigger. Figure 12 shows four squares in both the Poincaré and Klein models, one with material vertices, one with ideal vertices, and two with *hyperideal* vertices, that is vertices living beyond the boundary of \mathbb{H}^2 . Hyperideal points have a

dual relationship to planes within hyperbolic space [17, p. 71]. The largest hyperideal square has no material points and only four ideal points on the boundary of \mathbb{H}^2 . We will shortly see similar degenerate edges in some of our three-dimensional honeycombs.

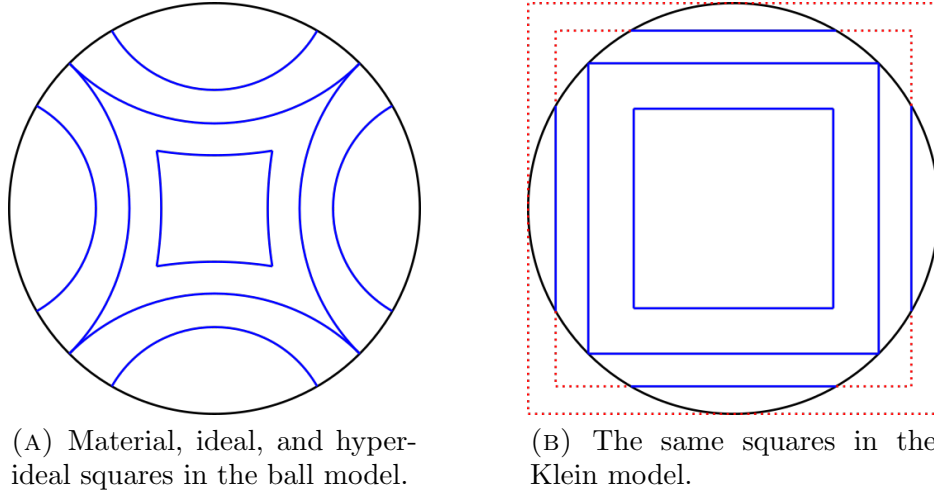


FIGURE 12. Hyperidealization. Hyperideal vertices of the squares can be seen in the Klein model.

Continuing our $\{4, 3, r\}$ sequence in three-dimensions, the $\{4, 3, 7\}$ honeycomb has cubical cells with hyperideal vertices. As in two dimensions, the vertices no longer exist, even on the boundary of \mathbb{H}^3 . Instead the cell has eight *legs* that end in triangular patches of the boundary of \mathbb{H}^3 . See Figure 13. In general, following Equation 1, a honeycomb $\{p, q, r\}$ has hyperideal vertices when $(r - 2)(q - 2) > 4$.






FIGURE 13. Three layers of cubes in the $\{4, 3, 7\}$ honeycomb. In these images, the corners of the cubes are cut off by the boundary of \mathbb{H}^3 , and you can see into the interiors of the cubes through the resulting holes.

When our honeycomb has hyperideal vertices, the geometry inside \mathbb{H}^3 becomes less interesting in a sense, because the edges never meet each other. It becomes harder to

understand what is going on by looking only at the edges going through the space, as we did in Figure 4. Instead, we can look at the pattern we get on the boundary of \mathbb{H}^3 from the hyperideal cells. Figure 14a shows one cube of the $\{4, 3, 7\}$ honeycomb. It intersects the boundary in eight triangles at the ends of its legs, which are drawn on the surface of the sphere in red. Also drawn are the intersections with the boundary of all other cubes in the honeycomb. Figure 14b shows this boundary pattern, stereographically projected from the sphere to the plane. Equivalently, we can think of this as being the boundary pattern of the $\{4, 3, 7\}$ honeycomb as drawn in the upper half space model of \mathbb{H}^3 . The triangular ends of the central cube are again shaded red, with progressively darker shades of blue for triangular ends of cubes farther away from the central cube. For more details on the colouring scheme used here, see Section 9.5.



(A) One cube with the induced pattern on the boundary of \mathbb{H}^3 .



(B) Stereographic projection of the boundary pattern.

FIGURE 14. The boundary pattern for the $\{4, 3, 7\}$ honeycomb. The eight *legs* of one cell are highlighted red. Compare with Figure 16.

Edges with ideal or hyperideal vertices usually intersect the boundary at two ideal points. However, as $r \rightarrow \infty$, less and less of each edge intersects \mathbb{H}^3 . Figure 15 shows a progression of $\{4, 3, r\}$ honeycombs as r becomes large. We can see from this picture where the endpoints of the edges are – for example two nearby white triangles have an edge between their closest vertices (compare with Figure 14). In the limiting $\{p, q, \infty\}$ honeycombs, each edge has only one ideal point, the rest of the edge having become hyperideal. In this sense, the edges of a $\{p, q, \infty\}$ honeycomb are like the edges of the largest square of Figure 12.

6.2. Hyperideal cells. As with ideal vertices, we can tell if a tiling will have hyperideal vertices by looking at the dual tiling and asking what kind of geometry the cells have. In the case of the $\{4, 3, 7\}$ honeycomb, the dual honeycomb has cells with Schläfli symbol $\{7, 3\}$, which is a tiling of \mathbb{H}^2 with three heptagons around each vertex. We can see many copies of the dual $\{3, 7\}$ tiling in Figure 14b, one for each hyperideal vertex: the






FIGURE 15. $\{4, 3, r\}$ honeycombs as $r \rightarrow \infty$. Compare with the $\{4, 3, 7\}$ honeycomb in Figure 14b. In the limit, the edges become hyperideal except for a single ideal point, located where the apparent $\{3, \infty\}$ tilings kiss.

hyperideal cells of the $\{7, 3, 4\}$ honeycomb are centered on the hyperideal vertices of the $\{4, 3, 7\}$ honeycomb, with a face corresponding to each edge of the $\{4, 3, 7\}$ honeycomb.

Hyperideal cells have a new feature: each such cell meets the boundary of \mathbb{H}^3 in a circular region we call the *head* of the cell. In general, a honeycomb $\{p, q, r\}$ has hyperideal cells when $(p - 2)(q - 2) > 4$.



(A) One $\{7, 3\}$ cell with the induced pattern on the boundary of \mathbb{H}^3 .



(B) Stereographic projection of the boundary pattern.

FIGURE 16. The boundary pattern for the $\{7, 3, 4\}$ honeycomb. The single *head* of one cell is highlighted red. Compare with Figure 14.

6.3. Ideal vertices and ideal cells. A honeycomb that has both ideal vertices and ideal cells cannot be represented in either a vertex-centered or a cell-centered picture. Instead, the most symmetrical pictures we can draw in the Poincaré ball model are edge-centered and face-centered. See for example Figure 17, showing a cell from the

$\{3, 6, 3\}$ honeycomb and a cell from the dual honeycomb, another copy of $\{3, 6, 3\}$. The blue cell is edge centered, with an edge of the cell going vertically through the center of the ball. The white cell is face centered, with a triangular face of the cell on the horizontal plane through the center of the ball. One ideal vertex of each cell is incident with the ideal center of its dual cell.




FIGURE 17. A cell from the $\{3, 6, 3\}$ honeycomb, overlaid with a dual cell.

6.4. Hyperideal vertices and hyperideal cells. Similarly, honeycombs with hyperideal vertices and hyperideal cells cannot be represented in either vertex-centered or cell-centered pictures. For example, Figure 18 shows a cell from the $\{3, 7, 3\}$ honeycomb and a cell from the dual honeycomb, another copy of $\{3, 7, 3\}$. As in Figure 17, the blue cell is edge centered, with an edge of the cell going vertically through the center of the ball. The white cell is face centered, with a hyperideal triangular face of the cell on the horizontal plane through the center of the ball. One hyperideal vertex of each cell is incident with the hyperideal center of the dual cell, visible as a leg of the cell protruding through the dual cell's head.




FIGURE 18. A cell from the $\{3, 7, 3\}$, overlaid with a dual cell.

7. THE SCHLÄFLI CUBE

We can now draw pictures of every Schläfli honeycomb. Each Schläfli honeycomb corresponds to an integral point of the infinite cube $\{(p, q, r) \mid 3 \leq p, q, r \leq \infty\}$. In this section we draw tables showing the honeycombs on the six faces of this cube, and on the two surfaces in Figure 11 corresponding to the honeycombs with ideal vertices and ideal cells. High resolution images of a selection of honeycombs throughout the cube are available at hyperbolichoneycombs.org. Section 9.2 describes the construction in more detail, and provides some general observations about the pattern on the boundary for arbitrary p , q , and r .

For honeycombs with neither hyperideal vertices nor hyperideal cells, we draw the in-space view (using POV-Ray [10]) – a view of the edges of the honeycomb as seen from within the space, assuming that light rays travel along geodesics. These are the fifteen hyperbolic honeycombs listed by Schlegel and Coxeter [2], the one euclidean honeycomb $\{4, 3, 4\}$, and the six spherical honeycombs corresponding to the regular four-dimensional polychora. For all other honeycombs, we draw the boundary picture, as described in Section 6.

A number of interesting patterns appear as the terms of the Schläfli symbol approach infinity. As $p \rightarrow \infty$, cell heads become tangent to each other. As $q \rightarrow \infty$, cell heads become tangent to hyperideal vertices. As $r \rightarrow \infty$, hyperideal vertices become tangent to each other. Only for $\{\infty, \infty, \infty\}$ do all of these tangencies exist simultaneously.

Once we started drawing pictures of hyperideal honeycombs, we began to see them in a number of places, especially honeycombs with hyperideal cells and material vertices. Many of these correspondences relate to Kleinian groups, fractals, and Apollonian gaskets. *Indra's Pearls*, by David Mumford, Caroline Series, and David Wright, shows an image of a Kleinian group, which turns out to be the $\{3, 3, \infty\}$ honeycomb [7, p. 200].

The cover of *Multifractals and 1/f Noise: Wild Self-Affinity in Physics* by Benoit B. Mandelbrot shows a fractal known as the “The Pharaoh’s Breastplate” [6]. This is an image of the $\{\infty, 3, 4\}$ honeycomb. Beautiful images of Schmidt arrangements in the paper *Visualising the arithmetic of imaginary quadratic fields* by Katherine E. Stange are intimately connected to ideal honeycombs such as the $\{4, 4, 4\}$ [16, p. 8]. Random groups are related to Coxeter groups of honeycombs of the form $\{p, 3, 3\}$ [1].

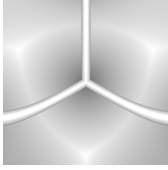

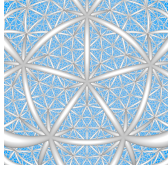
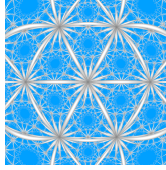
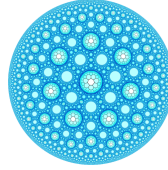
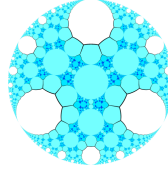
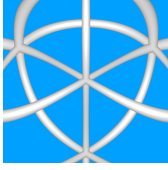
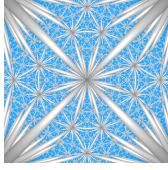
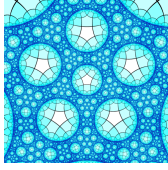
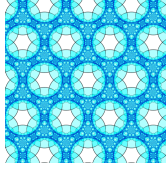
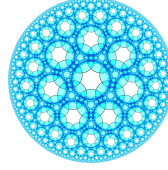
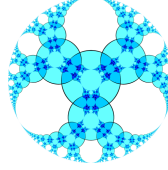
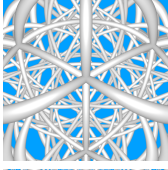
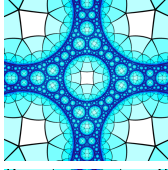
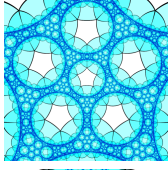
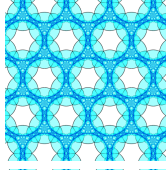
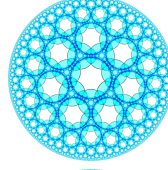
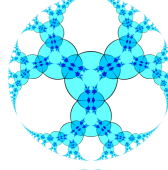
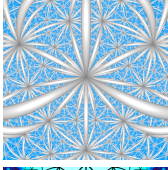
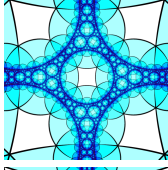
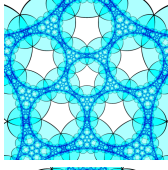
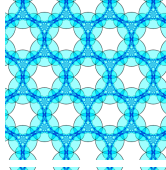
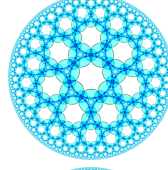
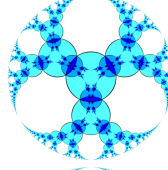
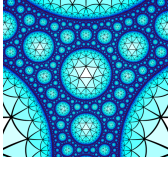
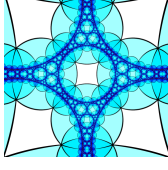
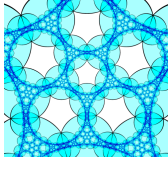
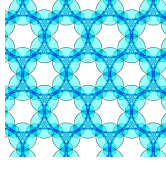
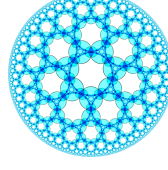
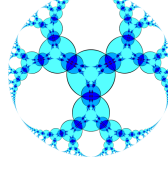
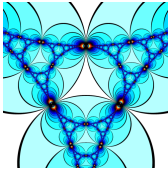
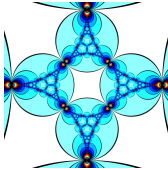
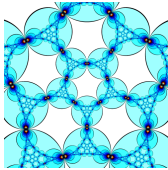
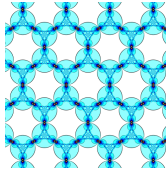
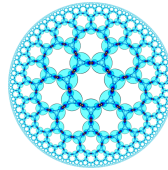
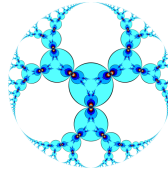
$r \backslash q$	3	4	5	6	7	...	∞
3						...	
4						...	
5						...	
6						...	
7						...	
...
∞						...	

TABLE 1. The $\{3, q, r\}$ honeycombs.

$r \backslash p$	3	4	5	6	7	\dots	∞
3						\dots	
4						\dots	
5						\dots	
6						\dots	
7						\dots	
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
∞						\dots	

TABLE 2. The $\{p, 3, r\}$ honeycombs.

$q \backslash p$	3	4	5	6	7	\dots	∞
3						\dots	
4						\dots	
5						\dots	
6						\dots	
7						\dots	
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
∞						\dots	

TABLE 3. The $\{p, q, 3\}$ honeycombs.

$r \backslash q$	3	4	5	6	7	\dots	∞
3						\dots	
4						\dots	
5						\dots	
6						\dots	
7						\dots	
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
∞						\dots	

TABLE 4. The $\{\infty, q, r\}$ honeycombs.

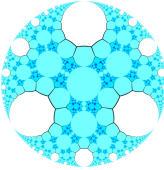
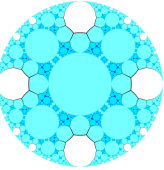
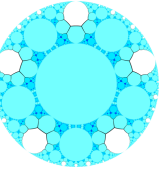
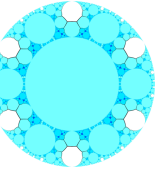
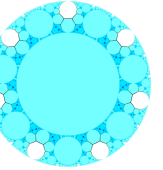
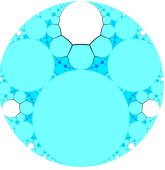
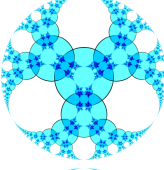
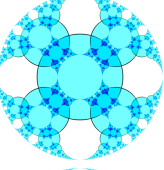
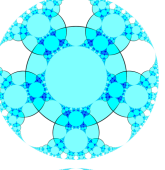
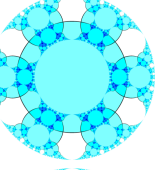
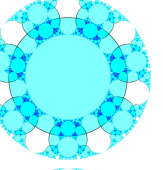
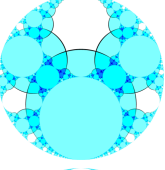
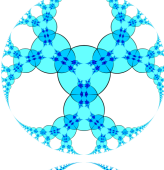
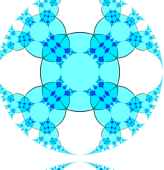
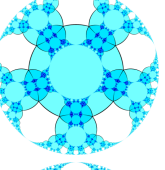
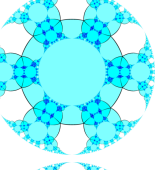
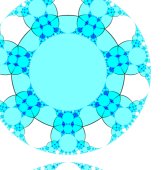
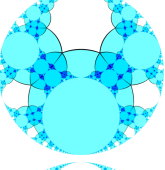
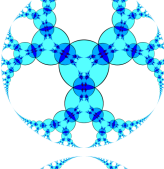
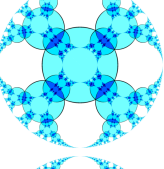
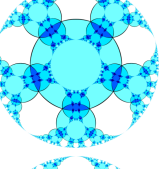
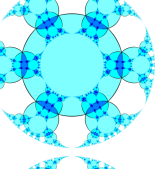
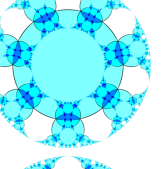
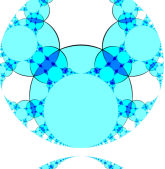
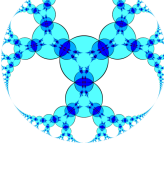
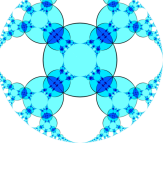
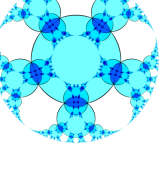
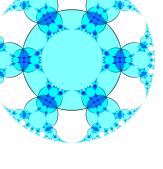
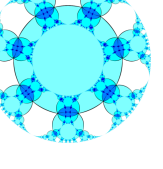
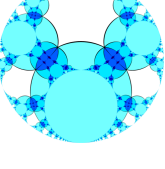
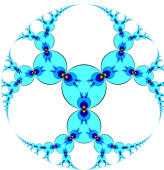
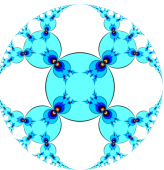
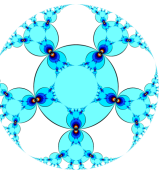
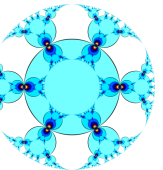
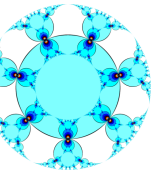
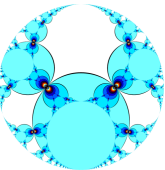
$r \backslash p$	3	4	5	6	7	\dots	∞
3						\dots	
4						\dots	
5						\dots	
6						\dots	
7						\dots	
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
∞						\dots	

TABLE 5. The $\{p, \infty, r\}$ honeycombs.

$q \backslash p$	3	4	5	6	7	\dots	∞
3						\dots	
4						\dots	
5						\dots	
6						\dots	
7						\dots	
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
∞						\dots	

TABLE 6. The $\{p, q, \infty\}$ honeycombs.

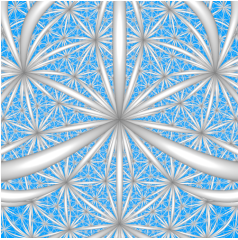
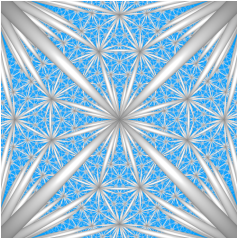
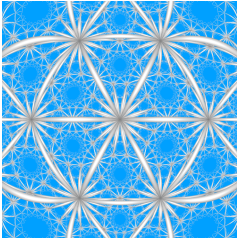
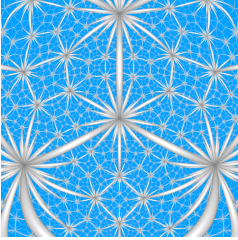
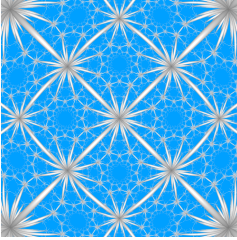
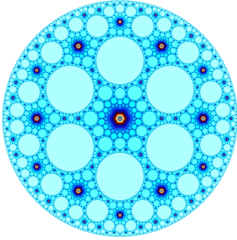
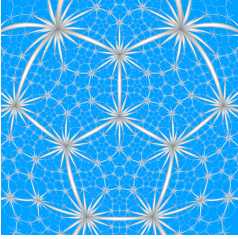
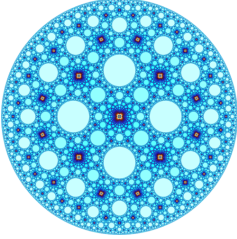
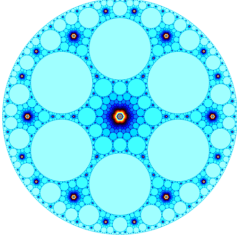
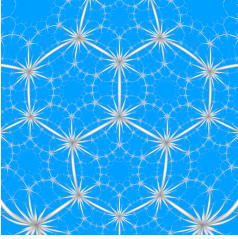
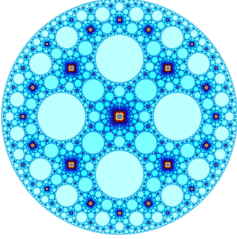
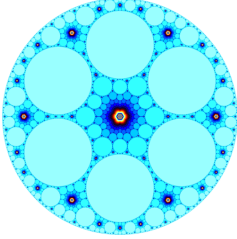
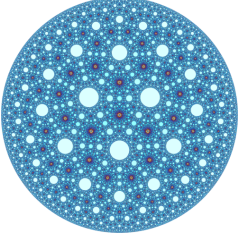
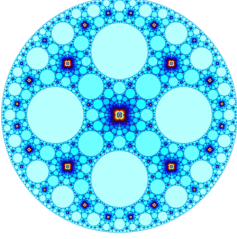
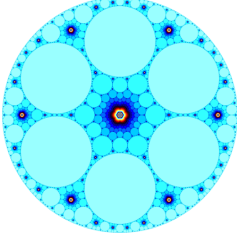
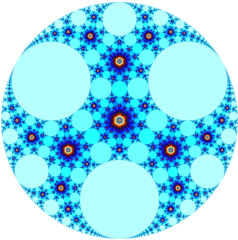
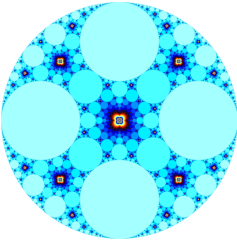
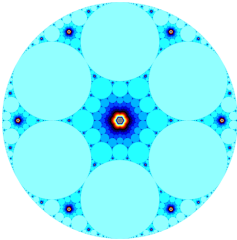
$p \setminus \{q, r\}$	$\{3, 5\}$	$\{4, 4\}$	$\{6, 3\}$
3			
4			
5			
6			
7			
\vdots	\vdots	\vdots	\vdots
∞			

TABLE 7. Honeycombs with ideal vertices.

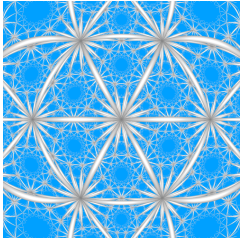
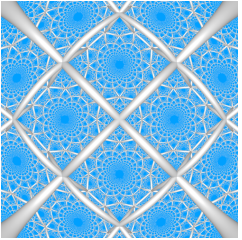
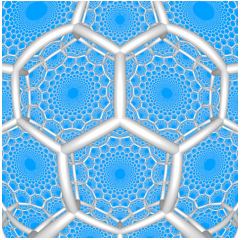
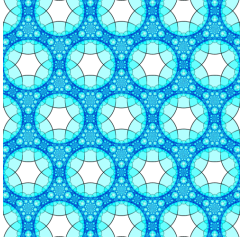
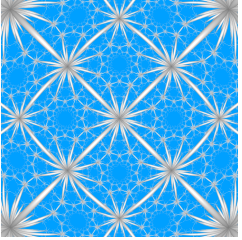
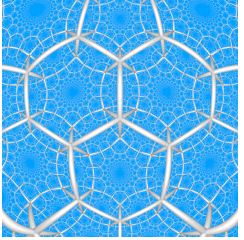
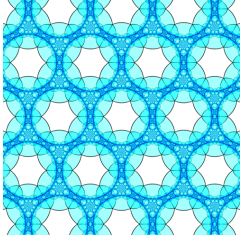
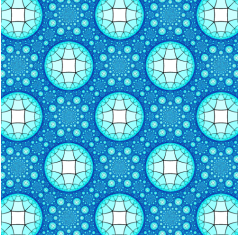
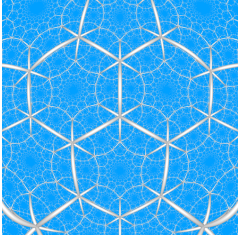
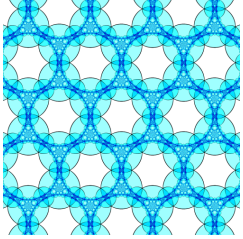
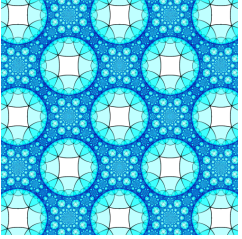
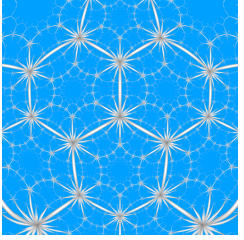
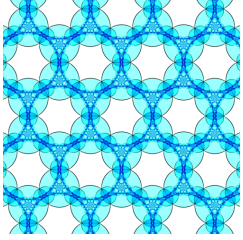
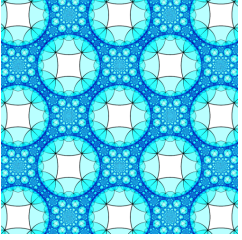
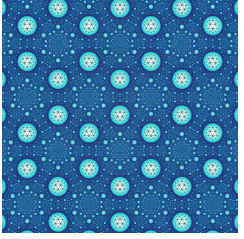
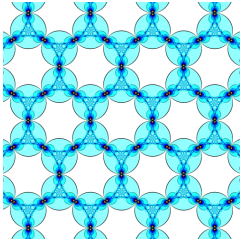
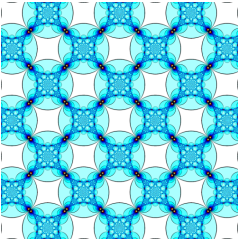
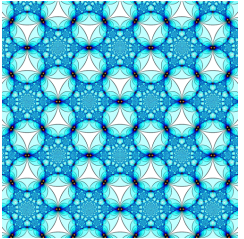
$r \setminus \{p, q\}$	$\{3, 6\}$	$\{4, 4\}$	$\{6, 3\}$
3			
4			
5			
6			
7			
\vdots	\vdots	\vdots	\vdots
∞			

TABLE 8. Honeycombs with ideal cells.

8. ARTWORK

8.1. 3D printed sculptures. A hyperbolic honeycomb represented using only edges as in Figure 5c is difficult to interpret from a two-dimensional picture, but a 3D printed sculpture is much more useful. See Figures 19 through 22. Also see Sections 9.3 and 9.4 for technical details on designing these sculptures.




FIGURE 19. The dual $\{4, 3, 5\}$ and $\{5, 3, 4\}$ honeycombs.




FIGURE 20. Two views of the self-dual $\{3, 5, 3\}$ honeycomb.




FIGURE 21. The dual $\{3, 3, 6\}$ and $\{6, 3, 3\}$ honeycombs.




FIGURE 22. Two views of the self-dual $\{4, 4, 4\}$ honeycomb. Only half of the Poincaré ball is rendered in this print.

8.2. Two-dimensional images. As we saw in Section 7, boundary patterns of hyperbolic honeycombs provide a huge collection of interesting and strikingly varied imagery. There are many other ways to show these honeycombs, giving yet more images.

One way we can alter the images is by applying hyperbolic isometries. Figure 23 shows part of the $\{4, 6, 4\}$ honeycomb, drawn on the boundary of the upper half space model of \mathbb{H}^3 , as in Section 7. Here however, we apply an isometry so that the head of the white cell takes up the whole of the lower half plane. Figure 24a shows a variation






FIGURE 23. Part of the boundary of the $\{4, 6, 4\}$ honeycomb.



(A) The $\{7, 3, \infty\}$ honeycomb, with a point at which two legs of a hyperideal cell meet sent to infinity. Compare with Table 6.



(B) The $\{4, 3, 7\}$ honeycomb, oriented so that the image is self-similar.

FIGURE 24. Further variations, applying isometries to the boundary image.

that results in a pattern with translational symmetry. Figure 24b shows the $\{4, 3, 7\}$ honeycomb moved so that the image is self-similar after a uniform scaling of

$$e^{2 \operatorname{arccosh}(\sqrt{2} \cos \frac{\pi}{7})} \approx 4.25917.$$

If we look down on the upper half space model, with the boundary of \mathbb{H}^3 being the page, then for this honeycomb we are looking down an infinite tower of cubes. The isometry that moves one cube in the tower to the next corresponds to scaling the boundary pattern, and so the boundary pattern is self-similar. The scaling factor may be calculated from the distance between faces in this tower of cubes, which is twice the inradius of the honeycomb. The inradius can be calculated using a formula due to Coxeter [2, p. 158].




FIGURE 25. Hyperbolic Catacombs.

Figure 25 is based on the $\{3, 7, 3\}$ honeycomb. The floor of the “catacombs” is the boundary of the $\{3, 7, 3\}$ honeycomb as shown in Tables 1 and 3. The image as a whole shows a view of the upper half space model of \mathbb{H}^3 , with only five cells of the honeycomb drawn: The ceiling of the catacombs is one cell, the other four are the family of red “creatures”.

9. IMPLEMENTATION DETAILS

Code for generating the images and models in this paper is available at [9].

9.1. Honeycomb construction. We construct general $\{p, q, r\}$ honeycombs in the conformal models: stereographic projection of S^3 for spherical space, ordinary euclidean space, and the upper half space model for hyperbolic space. In these models, geodesic surfaces are spheres (or planes, i.e. spheres of infinite radius), and we calculate the four surfaces for the fundamental simplex of the honeycomb. With the fundamental simplex defined, a honeycomb is built by recursive reflections (sphere inversions, see [8, pp. 124-126]) in the surfaces of the simplex. A robust sphere inversion function is important, which can handle all cases of spheres or planes reflecting to spheres or planes.


We use the four mirrors differently, depending on the need. For the 3D prints and POV-Ray renderings, we start with some honeycomb element, say an edge or a face, and reflect it around to build up the honeycomb, using some size threshold to stop the recursion. For the upper half space boundary images, we do the reverse. For every pixel in the image, we reflect it in the simplex mirrors until it is within the fundamental simplex. We then colour the pixel based on where it ends up in the simplex, and how many reflections have taken place.

9.2. Simplex construction. Calculating the geometry of the fundamental simplex relies on an observation about hyperideal honeycombs, namely that the upper half space pattern of a $\{p, q, r\}$ honeycomb visually looks like a set of solid disks (cell heads) and $\{q, r\}$ disks (hyperideal vertices) arranged in a *meta- $\{p, q\}$ -tiling* pattern. See Figure 26a. In the meta-tiling, the solid disks live at meta-tile centers and the $\{q, r\}$ disks live at meta-tile vertices. The meta-tiling may be spherical (as in the first three columns of Table 1), euclidean (as in Table 8), or hyperbolic (as in the right two columns of Table 1).

The first three of the simplex mirrors (drawn in red) reflect the simplex to another simplex within a cell of the honeycomb, and the last of the mirrors (drawn in orange) reflects the simplex to a simplex of a different cell. On the plane at infinity, the first three mirrors are exactly those of the meta-tiling. The mirror that reflects across cells is centered on a meta-cell, and its radius can be calculated using trigonometry so that r cells surround each edge.

Figure 26b shows the fundamental simplex of the $\{6, 5, 4\}$ honeycomb in the upper half space model. This simplex has two material vertices and two hyperideal vertices. The six edges of the simplex are drawn in black, with constant hyperbolic thickness. The simplex edges connected to a hyperideal vertex do not meet within the space, so this simplex itself has two legs, one at the left of the image (inside the cell leg in the center of the boundary pattern), and one at the right (inside the cell head centered at the $z = 0$ plane's point-at-infinity). There is only one edge that connects to both hyperideal vertices, which is the vertical edge at the left. The other end of this edge is not visible, since it meets the boundary at the $z = 0$ plane's point-at-infinity. There is also only one edge connected to both material vertices, in the central part of the image. The remaining four edges have one material and one hyperideal endpoint. The four simplex faces are rendered in semitransparent gray. Two are portions of planes and two are portions of spheres. All are geodesic and orthogonal to the boundary.

Infinite terms in the Schläfli symbols provide challenges in the same way that they do for two-dimensional tilings. For example, since an ∞ -gon does not have a material



(A) The boundary pattern. The meta- $\{6, 5\}$ -tiling, also generated by the red mirrors, is overlaid in black.

(B) A simplex of the $\{6, 5, 4\}$ honeycomb in the upper half space model.

FIGURE 26. The $\{6, 5, 4\}$ honeycomb. The intersection of the first three simplex mirror planes with the boundary plane are drawn red. The intersection of the mirror that reflects across cells is drawn orange.

center, we cannot draw a tile-centered $\{\infty, 3\}$ tiling. It is more natural to construct this tiling as the dual to a $\{3, \infty\}$ tiling. For three-dimensional honeycombs, we handle various combinations of infinite p , q , or r similarly, and our orientation choices were made to keep the visual progressions in our tables as consistent as possible.

9.3. Edges as Dupin cyclides. Edges, the one dimensional elements of a honeycomb, are geodesics in the induced geometry. To realize honeycomb edges in our physical models, we must give them some thickness. A constant offset from the line of an edge forms a *Clifford surface*. From within hyperbolic space, thickened edges appear cylindrical, but they take on a banana-like shape when projected to the Poincaré ball model (see Figures 19, 20, 21b and 22) or upper half space model (see Figure 26b). These shapes are known as *Dupin cyclides* [12]. Dupin cyclides are inversions of cylinders, tori, or cones.

A cyclide is the envelope of a smooth one-parameter family of spheres. We construct meshes to represent our cyclides by calculating the centers and radii of regularly spaced spheres in the cyclide, and placing mesh vertices on the tangent circles where each sphere touches the cyclide – the curve on the sphere furthest from the cyclide’s core geodesic.

Consider a sphere with hyperbolic center $\mathbf{p}_{\mathbb{H}}$ in the Poincaré ball model, and hyperbolic radius $r_{\mathbb{H}}$. Because of the distortions of the Poincaré ball model, the euclidean center $\mathbf{p}_{\mathbb{E}}$ of this sphere is not generally the same as $\mathbf{p}_{\mathbb{H}}$. The exception is at the origin, where the hyperbolic and euclidean centers coincide, and the euclidean radius is

$r_0 = \tanh(r_{\mathbb{H}}/2)$. For all other spheres, take $d = \|\mathbf{p}_{\mathbb{H}}\|$, then a straightforward calculation shows the sphere has euclidean center

$$(3) \quad \mathbf{p}_{\mathbb{E}} = \frac{\mathbf{p}_{\mathbb{H}}(r_0^2 - 1)}{d^2 r_0^2 - 1}$$

and euclidean radius

$$(4) \quad r_{\mathbb{E}} = \frac{r_0(d^2 - 1)}{d^2 r_0^2 - 1}.$$

Note that the $\mathbf{p}_{\mathbb{E}}$ is closer to the origin than $\mathbf{p}_{\mathbb{H}}$. The variable r_0 gives us a parameter to control the thickness of edges in our 3D printed model, balancing strength against cost and visibility into the interior.

9.4. 3D printing issues. In this section, we discuss some difficulties in designing 3D prints of hyperbolic honeycombs.

In the Poincaré ball model, features become very small as we approach the boundary of the ball. The 3D printers we used can print wire-like features with a minimum diameter of around 1mm, a restriction we quickly run into. We manage this issue in two ways.

For models with material vertices, such as the $\{6, 3, 3\}$ honeycomb shown in Figure 21b, we print only the edges whose thickness remains above this threshold. This first stage of culling leaves “dangling” edges, which are connected to the rest of the model only at one end. A second culling removes the dangling edges, a process which can reveal yet more dangling edges. We repeat until every edge is connected at both ends.

For models with ideal vertices, such as the $\{4, 4, 4\}$ honeycomb shown in Figure 22, all edges approach a thickness of zero near the boundary of the ball. For these, we artificially adjust the edge thickness. In Figure 22 the edges have accurate thicknesses until they shrink to 1mm, but remain a constant 1mm thickness beyond that. For the $\{3, 3, 6\}$ honeycomb in Figure 21a we used edges of constant thickness in \mathbb{R}^3 .

In order to reduce costs, we also hollow out the interior of the cyclides, while keeping their wall thickness within a separate threshold required by the printers.

A particular challenge for these models is mesh unioning. We initially generated models in which each edge was a separate mesh. For some of the simpler honeycombs, in which a small number of edges meet at a vertex, we were able (with some tinkering) to use tools such as netfabb [15] and MeshLab [14] to merge the meshes into a single mesh suitable for printing. For models such as the $\{4, 4, 4\}$ honeycomb, so many meshes meet at each vertex that those software packages fail. We wrote a function using the AbFab3D [13] library that converts our meshes into voxels, performs unioning in voxel-space, then converts back to a mesh from the voxel representation. Computer memory is the limiting factor for the resolution we can achieve using this method.

9.5. Colouring schemes. In the images in Section 7, we colour the cells according to distance from a central cell in the cell adjacency graph of the honeycomb. The

immediate neighbours of the central cell all have the same colour, as do all of the cells at distance two, and so on. We cycle through a set of colours, and repeat if necessary.

To generate the set of colours, we employ a simple colouring scheme based on traversing the edges of the RGB colour cube. With the red, green, and blue colour components of colour normalized between zero and one, the components of the colour are all zero, corresponding to black, at one vertex of this cube. At the antipode, the components are all one, corresponding to white. Starting from the white vertex, we move to cyan, then blue, then black, then red, yellow, and back to white. The full colour path traverses a skew hexagon along the edges of the RGB colour cube.

The first three edges of this hexagon are cool colours and the final three edges are warm colours, producing somewhat of a rainbow effect across the entire range. We chose the rate of movement along the hexagon so that the images mostly visibly expose the cool range. This rate is different for each Schläfli symbol, in order to best show the features of each image. The rate is chosen so that most of the pixels in the image have colours in the first half of the hexagon path. We calculated the cell-adjacency depth for each pixel, and then a mean depth for the entire picture. We experimented with a number of different functions of this mean to control the rate of movement, and settled on a rate proportional to the square root of the mean, which we felt gave the best visual result.


There is a choice of six different skew hexagons for a colouring path starting and ending at the white vertex. Some of the other hexagons produced striking results, and we also experimented by starting at various points along the hexagon. Figure 27 shows an artwork taking advantage of both of these changes. It starts near the black vertex and moves along the hexagon described above in the opposite direction.

Thought not directly related to colouring, we note here that for the boundary images, we found that antialiasing made a huge difference in the quality of small features. For each pixel, we sample 16 nearby points (in a 4×4 grid) and average the colour of all of these points.

9.6. Drawing curves on the upper half space boundary. In Figures 14 through 15, many images in Section 7, and Figures 23 through 27, there is a subtlety in how we draw the cross section of our honeycombs on the boundary of upper half space. The faces of the cells, even if given a constant thickness within the interior of the \mathbb{H}^3 , have a cross section on the boundary along a line with zero thickness. In fact, any two distinct points on the boundary are infinitely far apart from each other.

However, for aesthetic reasons it is preferable to draw lines with thickness. Figure 23 has many disks, each of which appears to contain a $\{6, 4\}$ tiling of the Poincaré disk model of the hyperbolic plane. Each of these disks is associated to a hyperideal vertex of the honeycomb, where infinitely many cells meet. But these disks are not two-dimensional tilings in the usual sense. The apparent edges are intersections of faces with the boundary, and the apparent tiles are intersections of cells with the boundary. A line of apparent edges travels through an infinite number of the vertex disks, expanding towards the center and shrinking towards the boundary of each disk.

We chose to render the vertex disks as geometrically accurate two-dimensional tilings, with geodesic edges rendered as constant thickness “bananas” within the hyperbolic

FIGURE 27. The $\{\infty, \infty, \infty\}$ honeycomb.

metric in each apparent Poincaré disk model of the hyperbolic plane. The method of achieving this effect involves slightly altering the simplex mirror that reflects across cells. The new mirror coincides with the original mirror on the boundary circle of its containing vertex disk, but diverges from the original in the disk interior. Any points that lie within the original simplex, but outside of the altered simplex are colored black. In all of our boundary pictures, the radius of each banana is 0.025, measured in the hyperbolic metric of each apparent Poincaré disk.

10. FUTURE DIRECTIONS

In this paper we considered only regular honeycombs. We rendered a number of in-space views of archimedean honeycombs for Wikipedia, but have yet to explore hyper-ideal archimedean honeycombs. By *archimedean*, we mean honeycombs not transitive on all flags – for example, honeycombs with archimedean cells or honeycombs with more than one cell type. By relaxing the regularity constraint, there are many further possibilities for 3D prints, as well as new upper half space boundary images.

Schläfli symbols can be generalized further, to four-dimensional tilings and higher. Four-dimensional tilings are difficult to visualize. For example we have to cheat in Figure 4a to visualize a single hypercube – really we are only seeing the edges of the hypercube. The four-dimensional tiling $\{4, 3, 3, 4\}$, consisting of infinitely many hypercubes, seems to be out of reach of our techniques. Paradoxically, we may have better luck visualizing four-dimensional tilings with hyperideal vertices of cells – because just as in three-dimensions, we can see what is happening on the boundary of \mathbb{H}^4 , which is S^3 . Using stereographic projection again, we can visualize S^3 .

ACKNOWLEDGMENTS

We would like to thank Don Hatch and Nan Ma for early discussion about visualizing hyperideal honeycombs. The exciting thread that seeded our approach of boundary images is available online [5]. We would also like to thank Tom Ruen for useful discussions, and for suggesting the possibility of rendering $\{\infty, \infty, \infty\}$.

REFERENCES

- [1] Danny Calegari and Henry Wilton. 3-manifolds everywhere. arXiv:1404.7043, 2014.
- [2] H. S. M. Coxeter. Regular honeycombs in hyperbolic space. In *Proceedings of the International Congress of Mathematicians of 1954*, 1954.
- [3] H. S. M. Coxeter. *Regular Polytopes*. Dover Publications, 1973.
- [4] Jacob Dix. Unbroken composition. <http://www.mobilography.net/unbroken-composition/>, 2013. Accessed: 2014-01-21.
- [5] 4D Cubing Yahoo Group. Hyperbolic honeycomb $\{7, 3, 3\}$. <https://groups.yahoo.com/neo/groups/4D.Cubing/conversations/messages/2291>, 2012. Accessed: 2015-10-10.
- [6] Benoit B. Mandelbrot. Multifractals and $1/f$ noise, 1998.
- [7] David Mumford, Caroline Series, and David Wright. *Indra’s pearls: the vision of Felix Klein*. Cambridge University Press, 2002.
- [8] Tristan Needham. *Visual Complex Analysis*. Oxford University Press, USA, 1999.
- [9] Roice Nelson. Software for generating images and models of hyperbolic honeycombs. <https://github.com/roice3/Honeycombs>, 2015.
- [10] Persistence of Vision Pty. Ltd. (2004). Persistence of vision raytracer (version 3.7). <http://www.povray.org/download/>, 2015.
- [11] Saul Schleimer and Henry Segerman. Sculptures in S^3 . In Robert Bosch, Douglas McKenna, and Reza Sarhangi, editors, *Proceedings of Bridges 2012: Mathematics, Music, Art, Architecture, Culture*, pages 103–110. Tessellations Publishing, 2012. Available online at <http://archive.bridgesmathart.org/2012/bridges2012-103.html>.
- [12] Michael Schrott and Boris Odehnal. Ortho-circles of Dupin cyclides. *J. Geom. Graph*, 10:73–98, 2006.
- [13] Abfab3d – toolkit for 3d printing. <http://abfab3d.com/>, 2014.
- [14] Meshlab – system for editing 3d triangular meshes. <http://meshlab.sourceforge.net/>, 2015.
- [15] netfabb – software for 3d printing. <https://netfabb.azurewebsites.net>, 2014.
- [16] Katherine E. Stange. Visualising the arithmetic of quadratic imaginary fields. arXiv:1410.0417, 2014.
- [17] William P. Thurston. *Three-Dimensional Geometry and Topology, Vol. 1*. Princeton University Press, 1997.

AUSTIN, TX USA

E-mail address: roice3@gmail.com

DEPARTMENT OF MATHEMATICS, OKLAHOMA STATE UNIVERSITY, STILLWATER, OK USA

E-mail address: segerman@math.okstate.edu